

Improving Reliability of Word Similarity Evaluation by Redesigning Annotation Task and Performance Measure

Oded Avraham and Yoav Goldberg

Computer Science Department

Bar-Ilan University

Ramat-Gan, Israel

{oavraham1, yoav.goldberg}@gmail.com

Abstract

We suggest a new method for creating and using gold-standard datasets for word similarity evaluation. Our goal is to improve the reliability of the evaluation, and we do this by redesigning the annotation task to achieve higher inter-rater agreement, and by defining a performance measure which takes the reliability of each annotation decision in the dataset into account.

1 Introduction

Computing similarity between words is a fundamental challenge in natural language processing. Given a pair of words, a similarity model $\text{sim}(w_1, w_2)$ should assign a score that reflects the level of similarity between them, e.g.: $\text{sim}(\text{singer}, \text{musician}) = 0.83$. While many methods for computing sim exist (e.g., taking the cosine between vector embeddings derived by word2vec (Mikolov et al., 2013)), there are currently no reliable measures of quality for such models. In the past few years, word similarity models show a consistent improvement in performance when evaluated using the conventional evaluation methods and datasets. But are these evaluation measures really reliable indicators of the model quality? Lately, Hill et al (2015) claimed that the answer is no. They identified several problems with the existing datasets, and created a new dataset – SimLex-999 – which does not suffer from them. However, we argue that there are inherent problems with conventional datasets and the method of using them that were not addressed in SimLex-999. We list these problems, and suggest a new and more reliable way of evaluating similarity models. We then report initial experiments on a dataset of Hebrew nouns similarity that we created according to our proposed method.

2 Existing Methods and Datasets for Word Similarity Evaluation

Over the years, several datasets have been used for evaluating word similarity models. Popular ones include RG (Rubenstein and Goodenough, 1965), WordSim-353 (Finkelstein et al., 2001), WS-Sim (Agirre et al., 2009) and MEN (Bruni et al., 2012). Each of these datasets is a collection of word pairs together with their similarity scores as assigned by human annotators. A model is evaluated by assigning a similarity score to each pair, sorting the pairs according to their similarity, and calculating the correlation (Spearman’s ρ) with the human ranking. Hill et al (2015) had made a comprehensive review of these datasets, and pointed out some common shortcomings they have. The main shortcoming discussed by Hill et al is the handling of *associated but dissimilar* words, e.g. (*singer, microphone*): in datasets which contain such pairs (WordSim and MEN) they are usually ranked high, sometimes even above pairs of similar words. This causes an undesirable penalization of models that apply the correct behavior (i.e., always prefer similar pairs over associated dissimilar ones). Other datasets (WS-Sim and RG) do not contain pairs of associated words pairs at all. Their absence makes these datasets unable to evaluate the models’ ability to distinct between associated and similar words. Another shortcoming mentioned by Hill et al (2015) is *low inter-rater agreement* over the human assigned similarity scores, which might have been caused by unclear instructions for the annotation task. As a result, state-of-the-art models reach the agreement ceiling for most of the datasets, while a simple manual evaluation will suggest that these models are still inferior to humans. In order to solve these shortcomings, Hill et al (2015) developed a new dataset – Simlex-999 – in which the instructions

presented to the annotators emphasized the difference between the terms associated and similar, and managed to solve the discussed problems.

While SimLex-999 was definitely a step in the right direction, we argue that there are more fundamental problems which all conventional methods, including SimLex-999, suffer from. In what follows, we describe each one of these problems.

3 Problems with the Existing Datasets

Before diving in, we define some terms we are about to use. Hill et al (2015) used the terms *similar* and *associated but dissimilar*, which they didn't formally connect to fine-grained semantic relations. However, by inspecting the average score per relation, they found a clear preference for hyponym-hypernym pairs (e.g. the scores of the pairs (*cat*, *pet*) and (*winter*, *season*) are much higher than those of the cohyponyms pair (*cat*, *dog*) and the antonyms pair (*winter*, *summer*)). Referring hyponym-hypernym pairs as *similar* may imply that a good similarity model should prefer hyponym-hypernym pairs over pairs of other relations, which is not always true since the desirable behavior is task-dependent. Therefore, we will use a different terminology: we use the term *preferred-relation* to denote the relation which the model should prefer, and *unpreferred-relation* to denote any other relation.

The first problem is *the use of rating scales*. Since the level of similarity is a relative measure, we would expect the annotation task to ask the annotator for a ranking. But in most of the existing datasets, the annotators were asked to assign a numeric score to each pair (e.g. 0-7 in SimLex-999), and a ranking was derived based on these scores. This choice is probably due to the fact that a ranking of hundreds of pairs is an exhausting task for humans. However, using rating scales makes the annotations vulnerable to a variety of biases (Friedman and Amoo, 1999). Bruni et al (2012) addressed this problem by asking the annotators to rank each pair in comparison to 50 randomly selected pairs. This is a reasonable compromise, but it still results in a daunting annotation task, and makes the quality of the dataset depend on a random selection of comparisons.

The second problem is *rating different relations on the same scale*. In Simlex-999, the annotators were instructed to assign low scores to unpreferred-relation pairs, but the decision of *how*

low was still up to the annotator. While some of these pairs were assigned very low scores (e.g. $\text{sim}(\text{smart}, \text{dumb}) = 0.55$), others got significantly higher ones (e.g. $\text{sim}(\text{winter}, \text{summer}) = 2.38$). A difference of 1.8 similarity scores should not be underestimated – in other cases it testifies to a true superiority of one pair over another, e.g.: $\text{sim}(\text{cab}, \text{taxi}) = 9.2$, $\text{sim}(\text{cab}, \text{car}) = 7.42$. The situation where an arbitrary decision of the annotators affects the model score, impairs the reliability of the evaluation: a model shouldn't be punished for preferring (*smart*, *dumb*) over (*winter*, *summer*) or vice versa, since this comparison is just ill-defined.

The third problem is *rating different target-words on the same scale*. Even within preferred-relation pairs, there are ill-defined comparisons, e.g.: (*cat*, *pet*) vs. (*winter*, *season*). It's quite unnatural to compare between pairs that have different target-words, in contrast to pairs which share the target word, like (*cat*, *pet*) vs. *cat*, *animal*). Penalizing a model for preferring (*cat*, *pet*) over (*winter*, *season*) or vice versa impairs the evaluation reliability.

The fourth problem is that *the evaluation measure does not consider annotation decisions reliability*. The conventional method measures the model score by calculating Spearman correlation between the model ranking and the annotators average ranking. This method ignores an important information source: the reliability of each annotation decision, which can be determined by the agreement of the annotators on this decision. For example, consider a dataset containing the pairs (*singer*, *person*), (*singer*, *performer*) and (*singer*, *musician*). Now let's assume that in the average annotator ranking, (*singer*, *performer*) is ranked above (*singer*, *person*) after 90% of the annotators assigned it with a higher score, and (*singer*, *musician*) is ranked above (*singer*, *performer*) after 51% percent of the annotators assigned it with a higher score. Considering this, we would like the evaluation measure to severely punish a model which prefers (*singer*, *person*) over (*singer*, *performer*), but be almost indifferent to the model's decision over (*singer*, *performer*) vs. (*singer*, *musician*) because it seems that even humans cannot reliably tell which one is more similar. In the conventional datasets, no information on reliability of ratings is supplied except for the overall agreement, and each average rank has the same weight in the evaluation measure. The problem of relia-

bility is addressed by Luong et al (2013) which included many rare words in their dataset, and thus allowed an annotator to indicate “Don’t know” for a pair if they does not know one of the words. The problem with applying this approach as a more general reliability indicator is that the annotator confidence level is subjective and not absolute.

4 Proposed Improvements

We suggest the following four improvements for handling these problems.

(1) The annotation task will be an explicit ranking task. Similarly to Bruni et al (2012), each pair will be directly compared with a subset of the other pairs. Unlike Bruni et al, each pair will be compared with only a few carefully selected pairs, following the principles in (2) and (3).

(2) A dataset will be focused on a single preferred-relation type (we can create other datasets for tasks in which the preferred-relation is different), and only preferred-relation pairs will be presented to the annotators. We suggest to spare the annotators the effort of considering the *type* of the similarity between words, in order to let them concentrate on the *strength* of the similarity. Word pairs following unpreferred-relations will not be included in the annotation task but will still be a part of the dataset – we always add them to the bottom of the ranking. For example, an annotator will be asked to rate (*cab*, *car*) and (*cab*, *taxi*), but not (*cab*, *driver*) – which will be ranked last since it’s an unpreferred-relation pair.

(3) Any pair will be compared only with pairs sharing the same target word. We suggest to make the pairs ranking more reliable by splitting it into multiple target-based rankings, e.g.: (*cat*, *pet*) will be compared with (*cat*, *animal*), but not with (*winter*, *season*) which belongs to another ranking.

(4) The dataset will include a reliability indicator for each annotators decision, based on the agreement between annotators. The reliability indicator will be used in the evaluation measure: a model will be penalized more for making wrong predictions on reliable rankings than on unreliable ones.

4.1 A Concrete Dataset

In this section we describe the structure of a dataset which applies the above improvements. First, we need to define the preferred-relation (to apply improvement (2)). In what follows we use the hyponym-hypernym relation. The dataset is

	w_t	w_1	w_2	$R_{>}(w_1, w_2; w_t)$
P	singer	person	musician	0.1
P	singer	artist	person	0.8
P	singer	musician	performer	0.6
D	singer	musician	song	1.0
R	singer	musician	laptop	1.0

Table 1: Binary Comparisons for the target word *singer*. P: positive pair; D: distractor pair; R: random pair.

based on *target words*. For each target word we create a group of *candidate words*, which we refer to as the *target-group*. Each candidate word belongs to one of three categories: *positives* (related to the target, and the type of the relation is the preferred one), *distractors* (related to the target, but the type of the relation is not the preferred one), and *randoms* (not related to the target at all). For example, for the target word *singer*, the target group may include *musician*, *performer*, *person* and *artist* as positives, *dancer* and *song* as distractors, and *laptop* as random. For each target word, the human annotators will be asked to rank the positive candidates by their similarity to the target word (improvements (1) & (3)). For example, a possible ranking may be: *musician* > *performer* > *artist* > *person*. The annotators responses allow us to create the actual dataset, which consists of a collection of *binary comparisons*. A binary comparison is a value $R_{>}(w_1, w_2; w_t)$ indicating how likely it is to rank the pair (w_t, w_1) higher than (w_t, w_2) , where w_t is a target word and w_1, w_2 are two candidate words. By definition, $R_{>}(w_1, w_2; w_t) = 1 - R_{>}(w_2, w_1; w_t)$. For each target-group, the dataset will contain a binary comparison for any possible combination of two positive candidates, as well as for all the combinations in which the first candidate is positive and the second is negative (either distractor or random). When comparing two positive candidates w_{p_1}, w_{p_2} – the value of $R_{>}(w_{p_1}, w_{p_2}; w_t)$ is the portion of annotators who ranked (w_t, w_{p_1}) over (w_t, w_{p_2}) . When comparing a positive candidate w_p to a negative one w_n – the value of $R_{>}(w_p, w_n; w_t)$ is 1. This reflects the intuition that a good model should always rank preferred-relation pairs above other pairs. Notice that $R_{>}(w_1, w_2; w_t)$ is the reliability indicator for each of the dataset key answers, which will be used to apply improvement (4). For some example comparisons, see Table 1.

4.2 Scoring Function

Given a similarity function between words $sim(x, y)$ and a triplet (w_t, w_1, w_2) let $\delta = 1$ if $sim(w_t, w_1) > sim(w_t, w_2)$ and $\delta = -1$ otherwise. The score $s(w_t, w_1, w_2)$ of the triplet is then: $s(w_t, w_1, w_2) = \delta(2R_>(w_1, w_2; w_t) - 1)$. This score ranges between -1 and 1 , is positive if the model ranking agrees with more than 50% of the annotators, and is 1 if it agrees with all of them. The score of the entire dataset C is then:

$$\frac{\sum_{w_t, w_1, w_2 \in C} \max(s(w_t, w_1, w_2), 0)}{\sum_{w_t, w_1, w_2 \in C} |s(w_t, w_1, w_2)|}$$

The model score will be 0 if it makes the wrong decision (i.e. assign a higher score to w_1 while the majority of the annotators ranked w_2 higher, or vice versa) in every comparison. If it always makes the right decision, its score will be 1. Notice that the *size* of the majority also plays a role. When the model takes the wrong decision in a comparison, nothing is being added to the numerator. When it takes the right decision, the numerator increase will be larger as reliable as the key answer is, and so is the general score (the denominator does not depend on the model decisions).

It worth mentioning that a score can also be computed over a subset of C , as comparisons of specific type (positive-positive, positive-distractor, positive-random). This allows the user of the dataset to make a finer-grained analysis of the evaluation results: it can get the quality of the model in specific tasks (preferring similar words over less similar, over words from unpreferred-relation, and over random words) rather than just the general quality.

5 Experiments

We created two datasets following the proposal discussed above: one preferring the hyponym-hypernym relation, and the other the cohyponym relation. The datasets contain Hebrew nouns, but such datasets can be created for different languages and parts of speech – providing that the language has basic lexical resources. For our dataset, we used a dictionary, an encyclopedia and a thesaurus to create the hyponym-hypernym pairs, and databases of word association norms (Rubinsten et al., 2005) and categories norms (Henik and Kaplan, 1988) to create the distractors pairs and the cohyponyms pairs, respectively. The

hyponym-hypernym dataset is based on 75 target-groups, each contains 3-6 positive pairs, 2 distractor pairs and one random pair, which sums up to 476 pairs. The cohyponym dataset is based on 30 target-groups, each contains 4 positive pairs, 1-2 distractor pairs and one random pair, which sums up to 207 pairs. We used the target groups to create 4 questionnaires: 3 for the hyponym-hypernym relation (each contains 25 target-groups), and one for the cohyponyms relation. We asked human annotators to order the positive pairs of each target-group by the similarity between their words. In order to prevent the annotators from confusing between the different aspects of similarity, each annotator was requested to answer only one of the questionnaires, and the instructions for each questionnaire included an example question which demonstrates what the term “similarity” means in that questionnaire (as shown in Figure 1).

Each target-group was ranked by 18-20 annotators. We measured the average pairwise inter-rater agreement, and as done in (Hill et al., 2015) – we excluded any annotator which its agreement with the other was more than one standard deviation below that average (17.8 percent of the annotators were excluded). The agreement was quite high (0.646 and 0.659 for hyponym-hypernym and cohyponyms target-groups, respectively), especially considering that in contrast to other datasets – our annotation task did not include pairs that are “trivial” to rank (e.g. random pairs). Finally, we used the remaining annotators responses to create the binary comparisons collection. The hyponym-hypernym dataset includes 1063 comparisons, while the cohyponym dataset includes 538 comparisons. To measure the gap between a human and a model performance on the dataset, we trained a word2vec (Mikolov et al., 2013) model¹ on the Hebrew Wikipedia. We used two methods of measuring: the first is the conventional way (Spearman correlation), and the second is the scoring method we described in the previous section, which we used to measure general and per-comparison-type scores. The results are presented in Table 2.

6 Conclusions

We presented a new method for creating and using datasets for word similarity, which improves

¹We used code.google.com/p/word2vec/ implementation, with window size of 2 and dimensionality of 200.

(A)	1st	2nd	3rd	4th	(B)	1st	2nd	3rd	4th
singer : artist	<input type="radio"/>	<input type="radio"/>	<input checked="" type="radio"/>	<input type="radio"/>	singer : dancer	<input checked="" type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>
singer : person	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input checked="" type="radio"/>	singer : chef	<input type="radio"/>	<input type="radio"/>	<input checked="" type="radio"/>	<input type="radio"/>
singer : musician	<input checked="" type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	singer : official	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input checked="" type="radio"/>
singer : performer	<input type="radio"/>	<input checked="" type="radio"/>	<input type="radio"/>	<input type="radio"/>	singer : painter	<input type="radio"/>	<input checked="" type="radio"/>	<input type="radio"/>	<input type="radio"/>

Figure 1: The example rankings we supplied to the annotators as a part of the questionnaires instructions (translated from Hebrew). Example (A) appeared in the hyponym-hypernym questionnaires, while (B) appeared in the cohyponyms questionnaire.

	Hyp.	Cohyp.
Inter-rater agreement	0.646	0.659
w2v correlation	0.451	0.587
w2v score (all)	0.718	0.864
w2v score (positive)	0.763	0.822
w2v score (distractor)	0.625	0.833
w2v score (random)	0.864	0.967

Table 2: The hyponym-hypernym dataset agreement (0.646) compares favorably with the agreement for nouns pairs reported by Hill et al (2015) (0.612), and it is much higher than the correlation score of the word2vec model. Notice that useful insights can be gained from the per-comparison-type analysis, like the model’s difficulty to distinguish hyponym-hypernym pairs from other relations.

evaluation reliability by redesigning the annotation task and the performance measure. We created two datasets for Hebrew and showed a high inter-rater agreement. Finally, we showed that the dataset can be used for a finer-grained analysis of the model quality. A future work can be applying this method to other languages and relation types.

Acknowledgements

The work was supported by the Israeli Science Foundation (grant number 1555/15). We thank Omer Levy for useful discussions.

References

- Eneko Agirre, Enrique Alfonseca, Keith Hall, Jana Kravalova, Marius Paşca, and Aitor Soroa. 2009. A study on similarity and relatedness using distributional and wordnet-based approaches. In *Proceedings of Human Language Technologies: The 2009 Annual Conference of the North American Chapter of the Association for Computational Linguistics*, pages 19–27. Association for Computational Linguistics.
- Elia Bruni, Gemma Boleda, Marco Baroni, and Nam-Khanh Tran. 2012. Distributional semantics in technicolor. In *Proceedings of the 50th Annual Meet-*

ing of the Association for Computational Linguistics: Long Papers-Volume 1, pages 136–145. Association for Computational Linguistics.

- Lev Finkelstein, Evgeniy Gabrilovich, Yossi Matias, Ehud Rivlin, Zach Solan, Gadi Wolfman, and Eytan Ruppin. 2001. Placing search in context: The concept revisited. In *Proceedings of the 10th international conference on World Wide Web*, pages 406–414. ACM.

- Hershey H Friedman and Taiwo Amoo. 1999. Rating the rating scales. *Friedman, Hershey H. and Amoo, Taiwo (1999). "Rating the Rating Scales." Journal of Marketing Management*, Winter, pages 114–123.

- Avishai Henik and Limor Kaplan. 1988. Category content: Findings for categories in hebrew and a comparison to findings in the us. *Psychologia: Israel Journal of Psychology*.

- Felix Hill, Roi Reichart, and Anna Korhonen. 2015. Simlex-999: Evaluating semantic models with (genuine) similarity estimation. *Computational Linguistics*.

- Thang Luong, Richard Socher, and Christopher D Manning. 2013. Better word representations with recursive neural networks for morphology. In *CoNLL*, pages 104–113. Citeseer.

- Tomas Mikolov, Kai Chen, Greg Corrado, and Jeffrey Dean. 2013. Efficient estimation of word representations in vector space. *arXiv preprint arXiv:1301.3781*.

- Herbert Rubenstein and John B Goodenough. 1965. Contextual correlates of synonymy. *Communications of the ACM*, 8(10):627–633.

- O Rubinsten, D Anaki, A Henik, S Drori, and Y Faran. 2005. Free association norms in the hebrew language. *Word norms in Hebrew*, pages 17–34.